deliver the right content at the right time. Finally, the need to engage learners is the large win for the organization and operations layer to invest in Web-based learning.

Finally, one output of this effort will be to develop a handbook describing the steps for creating a virtual learning environment between two or more sites connected with broadband technology. Our end goal is to produce a project methodology that relates to technology planning, equipment selection, acquisition, installation, testing, and integration; formulates an instructional design plan focusing on learning objectives, and Web-based delivery in content development; develops user guidelines for instructors, learners, managers, production, and technical staff; documents problems encountered and their solutions; and develops case studies to demonstrate examples of how the learning environment may be used.

REFERENCES

Advanced Distributed Learning Initiative. (2004, August). *SCORM Overview*. Retrieved August 30, 2004, from http://www.adlnet.org/index.cfm?fuseaction=scormabt

Aviation Industry CBL Committee. (2004, August). *Welcome to the AICC Web Site*. Retrieved August 30, 2004, from http://www.aicc.org

Cisco Systems. (2004, August). *Cisco AVVID: Introduction*. Retrieved August 30, 2004, from http://www.cisco.com/en/US/netsol/ns340/ns19/networking_solutions_market_segment_solutions_home.html

Dictionary.com. (2004, August). *Acronyms & Abbreviations*. Retrieved August 30, 2004, from http://www.acronymfinder.com/af-query.asp?p=dict&String=exact&Acronym=VOI

Microsoft Corporation. (2004). *MSDN XML Standards Reference*. Retrieved August 30, 2004, from http://msdn.microsoft.com/library/default.asp?url=/library/en-us/xmlsdk/html/oriXMLStandards.asp

Chapter 3

The ADL Vision and Getting from Here to There

J. D. Fletcher
Institute for Defense Analyses

The Advanced Distributed Learning (ADL) initiative was undertaken by the Department of Defense (DoD) at the request of the White House Office of Science and Technology Policy and in cooperation with the other federal agencies. It is intended to provide a model for all federal agencies to use in making education, training, and performance-aiding readily accessible anytime, anywhere. It is being developed through intense, frequent collaboration among industry, government, and academic participants. Its specifications are being adopted across Europe, Asia, the Pacific Rim, and the Americas.

ADL is expected to deliver both training, which, as a means to an end, prepares individuals to perform specific tasks and jobs, and education, which prepares individuals for life and is an end in itself. The knowledge representations and user interactions underlying these instructional applications are effectively identical to those needed to assist users in decision making, planning, problem solving, maintaining and operating equipment, and other performance-aiding functions. For these reasons, along with the costs to be saved by providing such functions (e.g., Fletcher & Johnston, 2002), performance aiding as well as instruction is a significant objective of the ADL initiative (Dodds & Fletcher, 2004; Wisher & Fletcher, 2004). However, performance aiding is a substantial topic by itself, and space limits discussion of its possibilities, promise, and implications here.

The implications of the ADL initiative are not limited to federal agencies. It presents opportunities and challenges in many other settings—including classroom instruction, especially as it is organized for K–16 education. ADL is not antithetical to such instruction. Its anytime, anywhere instructional goals include classrooms as well as workplaces, conference rooms, job sites, and homes. Nonetheless, the wide access to instruction that ADL will provide to students outside traditional instructional institutions and venues challenges instructors, administrators, and policymakers, as well as K–12 students and their parents, to assume new and unaccustomed roles with new and unaccustomed responsibilities.

With or without ADL, fully accessible anytime, anywhere education (and training and performance-aiding) is likely to become a reality. If nothing else, ADL is a harbinger of future learning processes, capabilities, and opportunities. Sooner or later, our instructional institutions must deal with them. If they do so successfully, their efforts will benefit everyone concerned with human learning and performance.

THE ADL VISION

The ADL initiative is based on the view of future education, training, and performance-aiding illustrated in Fig. 3.1. As the figure suggests, this view, or "vision," keys on three main components: (1) a global information grid—currently the World Wide Web—shown as the cloud on the left side of the figure, which provides an infrastructure populated by reusable instructional objects; (2) a server, shown in the middle of the figure, which locates and then assembles instructional objects into education, training, and/or performance-aiding materials tailored to user needs; and (3) devices, shown on the right side of the figure, which may be carried or worn and serve as personal learning associates that deliver education, training, and performance-aiding to users anywhere, anytime.

From an instructional point of view, the critical element in this vision is the server. It will assemble material needed to support interactions with learners and users on demand and in real time. These interactions will be tailored to the needs, capabilities, intentions, and learning state of each individual or group of individuals. The long-term vision of ADL is to establish instruction and performance-aiding interactions between technology and (human) users that consist of goal-directed *conversations* tailored to each learner's or user's needs, skills, knowledge, abilities, and interests.

Such conversations must draw on effective instructional strategies, accurate representations of the user, and comprehensive representations of relevant subject matter. Generative systems of this sort have been the objective of research and development investment by the DoD since the mid-1960s

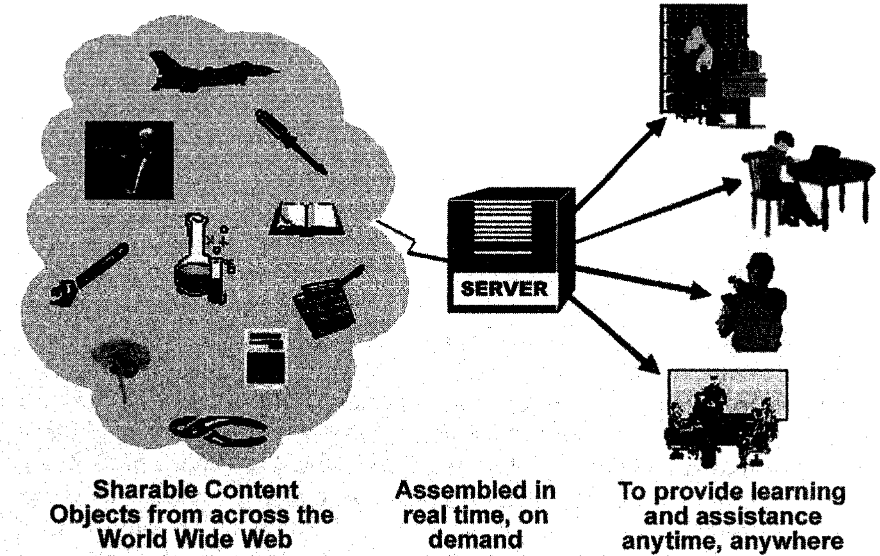

FIG. 3.1. An advanced distributed learning future.

(Carbonell, 1970; Fletcher, 1988; Fletcher & Rockway, 1986). They are the original goal of what today are called Intelligent Tutoring Systems (ITS). Generative capabilities were supported by the DoD, not so much as a clever application of artificial intelligence, but as a way to reduce the costs of instructional materials preparation by substituting the capital of technology and automation for human labor (Fletcher & Rockway, 1986). They remain essential to achieving the long-term goals of ADL.

The Basis for ADL

ADL goals arise from four main technical opportunities; advances in electronics, the pervasive accessibility of the World Wide Web, capabilities developed for Intelligent Tutoring Systems (ITS), and emerging specifications for reusable, sharable instructional objects.

Electronics. The first and most obvious opportunity keys on the rapid development of digital electronics spurred on by the operation of Moore's Law (Mann, 2000; Service, 1996). As readers may recall, Gordon Moore is a semiconductor pioneer and cofounder of the Intel Corporation. In 1965, Moore noted that engineers were doubling the number of electronic devices (basically transistors) on chips every year. In 1975, he revised his prediction to say that the doubling would occur every 2 years. If we split the difference and predict that it will occur every 18 months, our expectations fit reality quite closely.

A consequence of Moore's Law is that computers initially selling for $3,000 may cost about half that 18 months later. Another consequence is that the delivery devices shown on the right side of Fig. 3.1, will continue to decrease as much in physical size and cost as they will increase in functionality—not unlike today's cellular telephones, which themselves may provide ubiquitous platforms for ADL delivery.

World Wide Web. The second technological opportunity underlying ADL arises from the development and implementation of the global information grid, which currently takes the form of the World Wide Web. It was the Web that made feasible the DoD goal of accessible education, training, and performance-aiding available anytime and anywhere. Development of the Semantic Web will contribute substantially to achieving the long-term interactive instructional goals of ADL by providing the semantic linkages needed to create comprehensive representations of subject matter, expertise, learners, and users.

Intelligent Tutoring Systems. The third technological opportunity is presented by the generative capabilities of Intelligent Tutoring Systems (ITS). ADL development assumes that to be successful, its functionalities must be tailored to the specific needs, abilities, goals, and interests of the individual student or user (Dodds & Fletcher, 2004; Fletcher, 2003; Wisher & Fletcher, 2004). This tailoring, or individualization, is as critical for performance-aiding as it is for education and training. ADL functions are expected to provide what Brown, Burton, and DeKleer (1982) called articulate expertise. Not only must these functions provide helpful and relevant guidance, they must do so in a way that learners and users with varying levels of knowledge and skill can understand. ADL will, therefore, need to be "intelligent," building on capabilities that have been the developmental province of ITS.

At this point, it may be worth emphasizing the capabilities provided by "nonintelligent" computer-based instruction programs since the 1950s. They have been able to:

- Accommodate individual students' rate of progress, allowing as much or as little time as each student needs to reach instructional objectives;
- Tailor both the content and the sequence of instructional content to each student's needs;

- Make the instruction as easy or difficult, specific or abstract, applied or theoretical as necessary; and
- Adjust to students' most efficient learning styles (collaborative or individual, verbal or visual, etc.).

These capabilities have been described, discussed, and reviewed by Galanter (1959), Atkinson and Wilson (1969), Suppes and Morningstar (1972), Fletcher and Rockway (1986), and many others. To one degree or another, they have been implemented and available in computer-based instruction from its inception.

Intelligence in Intelligent Tutoring Systems is a different matter and more than a marketing term. When it was first introduced into computer-based instruction, it concerned quite specific goals. The distinction between ITS and other computer-based instruction was keyed to specific capabilities that were first targeted in the 1960s (Carbonell, 1970; Sleeman & Brown, 1982). Two of these defining capabilities are that ITS:

- Allow either the system or the student to ask open-ended questions and initiate instructional, "mixed-initiative" dialogue as needed or desired, and
- Generate instructional material and interactions on demand rather than require developers to foresee and prestore all materials and interactions needed to meet all possible eventualities.

Mixed-initiative dialogue requires a language for information retrieval, decision aiding, and instruction that is shared by both the system and the student/user. Natural language has been a frequent choice for this capability (e.g., Brown, Burton, & DeKleer, 1982; Graesser, Person, & Magliano, 1995), but the language of mathematics, mathematical logic, electronics, and other well-structured communication systems have also been used (Barr, Beard, & Atkinson, 1975; Psotka, Massey, & Mutter, 1988; Sleeman & Brown, 1982; Suppes, 1981; Woolf & Regian, 2000).

Generative capability requires the system to devise on demand—not draw from predicted and prestored formats—interactions with students. This capability involves not just generating problems tailored to each student's needs, but also coaching, hints, critiques of completed solutions, appropriate and effective teaching strategies, and, overall, the interactions and presentations needed for one-on-one tutorial instruction. These interactions must be generated from information primitives using an "instructional grammar" that is analogous to the deep structure grammar of linguistics.

Motivations for these two capabilities can be found in the perennial desire for cost containment accomplished by generating rather than anticipating and prestoring responses to all possible student states and actions. But they also arise from basic research on human learning, memory, perception, and cognition. In the 1960s–1970s, as documented by Neisser (1967), among others, the emphasis in basic research and understanding of human behavior shifted from the strict logical positivism of behavioral psychology, which focused on directly observable actions, to consideration of the internal, cognitive processes that were needed to explain empirically observed behavioral phenomena and are assumed to mediate and enable human learning.

The hallmark of this approach is the understanding that seeing, hearing, and remembering are all acts of *construction,* making more or less use of the limited stimulus information provided by our perceptual capabilities. Constructivist approaches are the subject of much current and relevant discussion in instructional research circles (e.g., Duffy & Jonassen, 1992; Tobias & Frase, 2000), but they are firmly grounded in the primordial foundations of scientific psychology. For instance, William James (1890/1950) stated his General Law of Perception as the following: "Whilst part of what we perceive comes through our senses from the object before us, another part (and it may be the larger part) always comes out of our mind" (p. 747).

In this sense, the generative capability sought by ADL and ITS developers is not merely something nice to have, but essential if we are to advance beyond the constraints of prescribed, prebranched, programmed learning, and ad hoc principles commonly used to design computer-based instruction. We need an interactive, generative capability if we are to deal successfully with the extent, variety, and mutability of human cognition.

As stated, the ADL vision is that training, education, and performance-aiding will take the form of human–computer conversations. This capability has been realized in systems that could converse in a formal language such as computer programming, e.g., BIP (Barr et al., 1975), or propositional calculus, e.g., EXCHECK (Suppes, 1981), or could base the conversation on determinate technical phenomena using clearly defined and well-understood terms, e.g., SOPHIE, (Brown et al., 1982). More recent research such as that discussed by Chipman (2003), Graesser, Gernsbacher, and Goldman (2003), and others suggests that significant natural language dialogue capabilities are achievable. Currently, however, they remain closer to the "bleeding edge" of technology than the mainstream. Although the ability to assemble reusable objects in real time and on demand into meaningful instructional or performance-aiding conversations has yet to be conclusively demonstrated, progress in the development of intelligent tutoring systems suggests that it is not an unreasonable goal.

Instructional Objects. Finally, ADL keys on technological opportunities offered by instructional objects. ADL development is presently focused on packaging instructional objects in anticipation of what has been called by Spohrer, Sumner, and Shum (1998) the "educational object economy." In such an economy, the emphasis in preparing materials for technology-based instruction (or performance-aiding) will shift from the current concern with developing instructional objects themselves to one of integrating already available objects into meaningful, relevant, and effective interactions.

The recent evolution of instructional objects has not escaped the attention of researchers. In assessing the educational value of objects, Roschelle and Kaput (1995) emphasized the ability to combine many kinds of interactive content in multiple display formats and obtain for education the benefits now being realized in business from the use of integrated office software. Roschelle et al. (1999) described software technologies underlying the development of five object-based education projects and reviewed their relative effectiveness. Gibbons, Nelson, and Richards (2000) reviewed in detail the nature and value of instructional objects for educational applications and concluded, along with Wiley (2000), that they are the technology of choice in supporting the evolution of technology-based instruction, because of their potential for reusability, adaptability, and scalability. Wiley (2000) provided the first scholarly, book-length treatment of instructional objects, and others may be expected to follow. However, beyond the economies, capabilities, and other benefits discussed by these researchers, is the possibility that the development of sharable instructional objects opens the door to genuinely generative instruction.

Instructional objects may then supply the primitives from which instructional interactions can be created on demand and in real time, and serve as reusable instructional components that reduce the costs of developing basic instructional materials—including graphics, simulations, and simulation scenarios (e.g., Towne, 1998, 2003). An economically viable, generative capability for instruction, in turn, depends on specifications for the development of sharable instructional objects. Developing these specifications has become a significant activity of the ADL initiative. The specifications are intended to separate objects from context-specific, run-time constraints, and proprietary systems so that the objects can be incorporated into other applications. They prescribe common interfaces and data interchange procedures so instructional objects can be aggregated into assemblies that guide and assist learners and users. Instructional objects so specified must be:

- *Accessible:* It should be possible to identify, locate, and access objects in one remote location and deliver them to many other locations when and where they are needed.
- *Interoperable:* Objects developed in one location with one set of tools or platform should be accessible and usable in other locations with dif-

ferent tools operating in different environments on different platforms.

- *Durable:* Objects should withstand technology changes (including version changes and upgrades) without redesign, reconfiguration, or recoding.
- *Reusable:* Objects should be sufficiently flexible to be independent of and used in multiple applications and contexts.

These criteria underlie SCORM, the Sharable Content Object Reference Model, which is an evolving specification for creating instructional objects that meet these criteria (Dodds, 2002; Dodds & Fletcher, 2004; Wisher & Fletcher, 2004). The SCORM specification is by no means the totality of ADL, but it provides a basis for populating the World Wide Web, or whatever form the global information grid takes in the future, with a ready supply of accessible, useable instructional objects. Once these objects exist, they must be identified, assessed, selected, and aggregated in real time and then handed to devices that deliver instruction and/or performance-aiding. As stated, this is the job of the server shown in the middle of Fig. 3.1. By importing "logic" or instructional strategy objects, as well as "content" objects, the server may also acquire expert tutoring capabilities.

This vision and view of instruction and performance-aiding conducted through human/computer conversations seem inevitable—sooner or later computer and instructional technology will produce personal learning associates with the functionalities envisioned by ADL. In its pursuit of instructional efficiency and knowledge accessibility, the ADL initiative seeks to bring it about sooner rather than later. Although difficult issues remain to be resolved, the most intransigent barriers to this future may be those arising from the current institutions, organizational structures, instructional practices, and administrative policies that are now vested in instruction conducted as lessons for students gathered together in one place and at one time. That is to say instruction conducted in classrooms, and not as tutorial conversations available to individuals anytime, anywhere. Although ADL technologies can be used as easily in classrooms as not, their impact will be to shift the emphasis in instructional practice, organizations, and approaches away from classrooms and onto individual users (Fletcher & Tobias, 2003).

Instructional Design

With SCORM, ADL has begun a process to populate the cloud on the left side of Fig. 3.1. The process assumes continued development of sharable instructional objects, continued development, expansion, and use of the World Wide Web, continued influence of Moore's Law on the cost and capabilities of electronics, and increasingly powerful, portable, and affordable electronic devices, which are depicted on the right side of Fig. 3.1. How might the server (in the middle of Fig. 3.1) operate to assemble instructional objects into effective instructional (and performance-aiding) interactions? How might it design these interactions?

At present, there is a gap between the one-of-a-kind, monolithic products of computer-based instruction development and Intelligent Tutoring Systems on one hand and the object-based development targeted by ADL on the other. The SCORM specifications are intended to make instructional components sharable and reusable by either computers and/or the human designers who assemble these components into instructional (and performance-aiding) interactions. These specifications attempt to bridge the gap between ADL goals and current instructional design by providing a foundation for object-based assembly that begins with relatively small, reusable learning resources, which are then aggregated to form units of instruction (or performance-aiding). By themselves, the objects may have no specific context. When they are combined with other objects, the resulting aggregation provides context and enables interactions with users in a sustained instructional or problem-solving conversation.

Working forward from SCORM we might well ask how its specifications and operating principles for instructional objects affect and shape the process of instructional design. SCORM has evolved through a series of versions, each intended to build on previous versions, rather than replace them. Its organizational structure may be described as a set of semi-independent functions. The SCORM specifications (currently available at www.adlnet.org) integrate and harmonize application details and requirements drawn from various standards and other specifications. They may be summarized as the following:

- *SCORM Content Aggregation Model* describes how to develop sharable instructional objects, package them for exchange from system to system, and describe them for search and discovery.
- *SCORM Run-Time Environment* describes the Learning Management System (or server) requirements needed to manage the ADL runtime environment, covering such matters as materials launch, communication between materials and server, and data model elements for sharing information about the learner's progress and needs.
- *SCORM Sequencing and Navigation* describes how SCORM materials may be sequenced through a set of learner-initiated or system-initiated navigation events. Branching and flow may be described by a predefined set of activities, determined at design time, or generated as needed on demand.

Although the SCORM Content Aggregation and Run–Time Environment specifications present functions that are essential to ADL, it is SCORM Sequencing and Navigation that most affects the process of instructional design. With "sequencing" we find software engineers discussing what instructional designers might call "branching," and with "navigation" we find them addressing issues that instructional designers might call "learner control." The SCORM specifications ensure that sequencing and navigation implemented in any conformant development environment using any conformant tools will operate successfully and as required by the developer(s) in any conformant run-time environment.

The SCORM Sequencing and Navigation specifications are based on "use cases" drawn from already delivered and currently used technology-based instruction. These examples and the Sequencing and Navigation specifications developed for them key on techniques that are the base of what instructional designers may recognize as Keller's Personalized System of Instruction (PSI) and Crowder's intrinsic programming.

Keller's PSI involves a process of breaking up a course of instruction into an ordered series of modules and then pretesting students for their mastery of each module's content before beginning work in it (Keller, 1968). Students who pass the pretest skip the module and proceed to a pretest for the next module in the series. Students who do not pass the pretest are required to complete the module and then be retested. They repeat this process until they pass the test, in practice one of several parallel tests, and only then proceed on to the next module.

Various studies found PSI to be effective. A meta-analysis of 75 empirical comparisons of PSI with standard classroom practices was reported by Kulik, Kulik, and Cohen (1979). They found that the PSI programs they reviewed raised final examination scores by about 0.50 standard deviations, roughly an increase in the performance of 50th percentile students to that of 69th percentile students. They also found that PSI produced less variation in achievement, higher student ratings, and fewer course withdrawals. Despite these favorable results, Keller (1985) grew pessimistic about the use of PSI because of the substantial amount of instructor time required to set up PSI courses and the lack of support (mostly in the form of release time) from administrators.

Crowder's intrinsic programming (e.g., 1959) is interesting because it and not Skinner's approach (e.g., 1954) is the one almost exclusively found in practice—even though Skinner's extrinsic programming is most frequently cited as the backbone of programmed instruction. Crowder's approach will be familiar to most instructional developers. An example is the following:

In the multiplication $3 \times 4 = 12$, the number 12 is called a ________

A: Factor	{Branch to remedial X1}
B: Quotient	{Branch to remedial X2}
C: Product	{Reinforce, go to the next item}
D: Power	{Branch to remedial X3}

In this item, the system, the computer instructor, assumes that a student responding "A" misunderstands the meaning of "Factor," or lacks an understanding of "Product," or both. The student is branched to instructional items intended to correct one and/or the other of these cognitive states and then is returned to this or a similar item to try again. A similar remedial approach is applied to responses of "B" and "D." A student responding with "C" is usually rewarded, "reinforced," with encouraging, positive feedback and then sent on to whatever item best continues his or her progress through the instruction, an action that by itself may constitute positive reinforcement.

These two approaches for sequencing between modules (PSI) and within modules (intrinsic programming) provide the basis for much computer-based instruction delivered today and, either consciously or not, for the use cases that the SCORM Sequencing and Navigation specifications were intended to accommodate. A comparison between computer-based instruction branching and SCORM sequencing is shown in Fig. 3.2. The typical computer-based instruction branching structure on the left would be implemented as shown on the right by using SCORM Sequencing and Navigation. Both sides of the figure depict typical Keller PSI branching, although it is not difficult to see how the same schemes might be extended to within-module Crowder intrinsic programming. Boxes A, B, and C may all be instructional objects (even if one or more of them is an assessment module), or they may be assembled from more granular objects by the server.

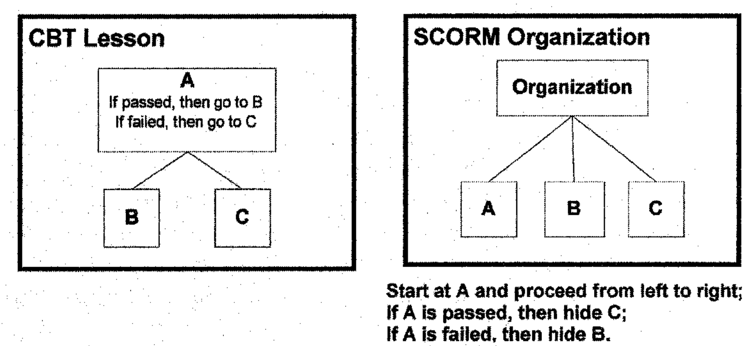

FIG. 3.2. Comparison of Computer-based Instruction Branching and SCORM Navigation.

The elements shown in Fig. 3.2 can be used as a basis for an entire course by repeating and extending them into a comprehensive tree structure that starts with a single entry point and branches out to as many modules as are needed. In this way, SCORM Sequencing and Navigation accommodate Keller's and Crowder's approaches and others as well. Learners proceed through the tree from module to module until they run out of modules and can proceed no further. At that point they are done with the course, unit, or lesson.[1]

The SCORM Sequencing and Navigation specifications assume that most learning design strategies can be represented as a tree of learning modules and that instructional designers will be able to map their instructional designs and strategies into this structure. SCORM allows access to global variables that can be used to communicate data from one module or object in the tree to another. It also allows modules to be invisible to learners thereby permitting use of specialized computations to implement more elaborate instructional approaches such as those involving Bayesian networks (e.g., Conati, Gertner, VanLehn, & Druzdzel, 1997; Corbett, Koedinger, & Anderson, 1997). These approaches should be able to function in the SCORM tree structure, determining if users are on the right track, shaping the scope and sequence of instruction they receive, and guiding their progress. Hidden Markov Models such as those used by Soller and Lesgold (2003) to facilitate collaboration among physically distributed participants may also be incorporated in a similar fashion to help learners collaborate in solving problems.

SCORM Sequencing and Navigation specifications are new and await empirical validation to determine the extent to which they support these and other approaches. The question naturally arises as to how much farther SCORM Sequencing and Navigation can be pushed toward the ADL vision in which instruction consists less of predefined lessons, testing, and screening and becomes more of a conversation between the learner/user and the technology.

Empirical demonstrations to answer this question may be showcased through "Designfests" that are analogs of the current ADL/SCORM "Plugfests." The ADL initiative uses Plugfests to assess and demonstrate how well servers, objects, and authoring tools conform to SCORM specifications and, in turn, how satisfactory the specifications are in meeting users' needs. In a similar fashion, Designfests will demonstrate how and to what extent SCORM specifications can be used to implement the instructional designs intentions of developers. Designfests will help identify gaps between what developers want to do and what is practicable under SCORM Sequencing and Navigation specification. They will also set priorities for further development of SCORM specifications.

[1]Fig. 3.2 is adapted from *SCORM Best Practices Guide for Content Developers* (2003), prepared by the Carnegie Mellon Learning Systems Architecture Lab to assist instructional developers. Readers who seek more information about implementing instructional strategies under SCORM Sequencing and Navigation are encouraged to consult the *Guide*.

TOWARD MORE ADAPTIVE, INTELLIGENT LEARNING SYSTEMS

Mainstream learning systems often rely on predetermined and fixed-path delivery of content. Such systems lack agility in adapting to learners' mastery states, and are thereby limited in their ability to tailor learning experiences to individual learners. As specified long ago (e.g., Fletcher, 1975), an adaptive, "intelligent" learning system needs an accurate model of the learner, a model of the knowledge domain, and a capability that can evaluate the differences between the two. It can then identify and/or devise, on demand and in real time, instructional strategies that produce desired instructional outcomes.

SCORM provides globally accessible records that can store the learner's degree of mastery. A hook was included in the records that permits them to reference externally defined competencies. As the learner is sequenced through the content objects, the learning system builds up a representation of the learner's mastery and progress. Records of this sort comprise a simple, accumulative model of the learner's level of competency in the area of instructional interest.

IMS Reusable Definition of Competency or Educational Objective (2002) adds to this capability by defining a taxonomy of competencies required to meet specific learning objectives. This taxonomy may be organized hierarchically to represent dependencies, supporting skills, or prerequisites. Each competency definition includes a text description of the competency and a unique identifier that may be referenced externally. The organization of a competency definition may represent specific skills or knowledge to be acquired for a specific task or subject domain. By referencing competency model identifiers, SCORM records can be used to compare the state of the learner with the generic IMS Global Learning Consortium competencies. This capability provides a generalizable, system-based means to perform knowledge and skills gap analyses leading to more sophisticated and adaptive strategies that use such information (Wiley, 2000).

As learning system specifications become more robust, they will also become more adaptive. Improved assessment methods and results are emerging that will continuously and unobtrusively extract information from instructional interactions and better represent the state of the learner. The strategies developed by learning systems will further be informed by learner profile information, which can "preload" the learner model with mastery information from outside sources, thereby reducing the need for additional testing to determine the learner's state. This process enhances the capabilities of technology-based instructional systems to bypass relevant content of already mastered material and concentrate on relevant material yet to be learned—a process that has long been advocated by researchers (e.g., Fletcher, 1975; Tobias, 1989).

Basically we seek an engineering of instruction (e.g., Woolf & Regian, 2000) with well-articulated principles for adjusting and modulating learning experiences. Such engineering would ensure that outcomes such as retention of skills and knowledge, application and transfer of learning, motivation to continue study, speed of response, accuracy of response, and so forth are reliably achieved by each learner to the maximum extent possible within the constraints imposed by instructional time and resources. This instructional engineering would automatically identify and devise learning strategies, and locate and assemble precisely appropriate objects into successive interactions with the learner. Each interaction would be tailored, on demand and in real time, to the outcome being sought, the learner's level of knowledge, skill, and style of learning, and the instructional strategy that was indicated by instructional principles. This is a significant challenge for instructional objects, Web-based services, and the state of the art in general, but current progress suggests that they may eventually rise to meet it, yielding technology that ensures reliable achievement of targeted instructional outcomes.

The Impact of Web and Web Services on this Evolution

One way the current and near-term capabilities of learning systems may evolve is through the Semantic Web, which will provide powerful new technologies for both knowledge representation and the ontologies needed to connect them (Berners-Lee, Hendler, & Lassila, 2001). These technologies will provide ways not only to relate but also to reason about information from widely different domains.

The Semantic Web is intended to imbue information available on the Web with sufficient meaning to significantly improve the cooperation between computers and human beings. Dealing with the semantic content of Web pages and information will enhance the process of discovery needed to access relevant information and objects from the Web. Access to this semantic content will key on the development, implementation, and use of ontologies, which make it possible to identify and expose semantic linkages between highly disparate bodies of information (Chandrasekaran, Josephson, & Benjamins, 1999).

Roughly, an ontology consists of a taxonomy and a set of inference rules that formally define operations and relations among the classes defined by the taxonomy. More specifically, ontologies consist of consensual, shared, formal descriptions that identify classes of objects, each member of which possesses all the qualities that all other members of the class have in common. The classes are organized in hierarchies, and classes of classes can be developed to any necessary depth. Relationships between a member of any class defined by an ontology can not only be quickly linked to many other classes and class members but the semantic quality that forms the link can also be exposed. Ontologies thereby identify semantic links between what may appear to be quite disparate classes and class members. Web services are being devised and implemented to identify and exploit these semantic linkages, and, in general, increase the "behavioral intelligence" of Web-based applications—as Bryson, Martin, McIlraith, and Stein (2002) have suggested.

These Web services are being built on top of existing and emerging Web standards, such as Hyper-Text Transfer Protocol (HTTP), Extensible Markup Language (XML), Universal Description, Discovery, and Integration (UDDI), and Simple Object Access Protocol (SOAP). In this way, emerging services are being made language, platform, and object model independent. They enable different applications running on different operating systems developed with different object models using different programming languages and programming environments to cooperate, communicate, and interoperate. They can express complex relationships using inference rules like those of intelligent tutoring systems to perform specific tasks such as profiling learners, representing their skills, knowledge, and abilities, linking these representations to instructional objects, and managing their progress toward instructional objectives and competencies.

If successful, the Semantic Web will integrate real-world knowledge and skills acquired through simulation, education, training, performance-aiding, and experience. It will provide a foundation for building more comprehensive and substantive models of subject matter domains and learners' levels of mastery than we now have and will combine them with more precise discovery of the instructional objects learners and other users need to develop desired human competencies. Building on the already available functionalities of intelligent tutoring systems, sharable objects, and existing standards, the Semantic Web and its services will contribute substantially to the next generation of learning environments.

Content Object Discovery and Retrieval

Given these considerations, it is not surprising to find that the development of Web services used to identify and retrieve contextually relevant instructional content is becoming a major topic. The success of Google and other Web search engines has demonstrated the value and utility of content discovery and whetted everyone's appetite for rapid, accurate search and retrieval. Presently, Google may be the single most important, effective, and widely used source of Web-based education. However, Google's location of content by text crawling, indexing, and retrieving everything that is remotely relevant to a search limits its use as a discovery system for focused content assembly. Its operation could be substantially improved if it were to cooperate with content and retrieve only what is intentionally prepared and published for discovery.

More precise identification of content objects is being addressed through the use of Uniform Resource Names (URN; http://www.faqs.org/rfcs/rfc1737.html), which serve as persistent, location-independent resource identifiers. The Corporation for National Research Initiatives (CNRI) has created a URN implementation called "The Handle System" (Kahn & Wilensky, 1995; http://www.handle.net/introduction.html). This system allows digital objects to obtain a unique identifier and to link each object to its location—wherever that might be—through the use of a Handle Resolution Service (similar to domain names resolving to Internet protocol addresses through the Domain Name System). CNRI hosts a global root server that can be queried to resolve requests.

Also, the Common Indexing Protocol (CIP; http://www.faqs.org/rfcs/rfc2651.html) allows the owner of content to create its index metadata while also allowing this indexing information to be shared among different servers, thereby enabling the development of new search and discovery services. New learning and performance-aiding specifications are emerging that permit the identification of skills, competencies, and knowledge so that logical relations among them that are relevant to specific but quite different communities of practices can be identified and then represented. As suggested earlier, not only will such logical relations be discovered, but the semantic nature of these relationships, insofar as they are reflected in metadata definitions, will be exposed.

These developments, among others, may produce Web services that provide accurate, precisely focused, and contextually correct discovery and retrieval of instructional objects on an easily scalable basis. Their combination of agility and accuracy enables considerable flexibility in dealing with the idiosyncratic prior knowledge elements and associations built up by individual users. They will allow instructional programs to continuously and unobtrusively assemble models of each user's state of knowledge, style of learning, and progress toward instructional objectives. These models will in turn support the precise tailoring of instructional interactions to each student that is a characteristic and unique strength of one-on-one tutoring—they will provide an Aristotle for every Alexander and a Mark Hopkins for the rest of us.

Where Might These Capabilities Take Us?

The emphasis on instructional technology brings us to revolutions in instruction. The first of these may have occurred with the development of written language about 7,000 years ago. It allowed the content of advanced ideas and teaching to transcend time and place. The second revolution in instruction began with the technology of books. Books made the content of high-quality instruction available anywhere and anytime, but also inexpensive and thereby accessible to many more people. A third revolution in instruction appears to be accompanying the introduction of computer technology. The capability of this technology for real-time adjustment of instructional content, sequence, scope, difficulty, and style to meet the needs of individuals suggests a third pervasive and significant revolution in instruction. It makes both the content *and* the interactions of high-quality instruction widely and inexpensively accessible—again anytime, anywhere.

Building on this possibility, ADL, SCORM, intelligent tutoring, and the Semantic Web in some combination may provide a foundation for generative education, training, and performance-aiding capabilities that are available anytime, anywhere. These developments can capitalize on the growth of electronic commerce and the World Wide Web. They can build on this worldwide, almost irresistible activity, accelerate it, and apply it to a full spectrum of education, training, and performance-aiding needs. But to realize all this promise, we must also learn to combine the software engineering features offered by SCORM with the best we have to offer in the form of instructional design.

The long-term, anytime, anywhere vision for ADL differs substantially from classroom learning and the many organizational structures we have in place to support it. But ADL is not at odds with classroom practice. Anytime, anywhere includes classrooms, and ADL capabilities are as accessible in classrooms as elsewhere. The instructional and performance-aiding, human–computer conversations that are the eventual goal of ADL will access the comprehensive spectrum of human knowledge becoming available from the World Wide Web and tailor it to the user's needs. These conversations will initially be designed to mimic those that are established by human tutors, but sooner or later this guiding metaphor must evolve and these conversations will take on forms, capabilities, and infrastructure of their own. The "Columbus Effect" will take over just as it did for wireless telegraph, horseless carriages, and a host of other technological innovations that led us into territory not envisioned in the original enabling metaphor.

At least three capabilities may evolve from the ADL teaching–learning environment:

- Less predefined sequencing—An instruction (or performance-aiding) conversation will presumably take whatever direction is needed by participants in the conversation. How to develop and provide a capability that allows sequences to adjust and evolve continually—perhaps a meta-sequencing capability—is a significant challenge for instructional designers. The notion of instructional design as a process of prespecifying and predefining a sequence of activities within a lesson module will need to evolve substantially if the ADL vision of a conversation sustained on an interaction-to-interaction basis is to be fully realized.
- More assessment and fewer tests—Assessment will become more continuous and unobtrusive as the capability for developing a model of the learner/user from interactions evolves. Such assessment may be accomplished by taking account of the learner's vocabulary, use of technical information, level of abstraction, clustering (chunking) of concepts, inferred hypothesis formation, and the like. These capabilities have yet to be fully explored and verified, but enough research on their application has been completed to suggest their promise for the continuous and unobtrusive assessment of user knowledge and abilities needed to tailor instruction and performance-aiding to their needs. Some explicit testing and explicit probing may still be used to assess learner progress efficiently. What sort of probes are needed, how they are to be implemented, and what principles will guide their psychometric properties is another challenge for instructional designers—a challenge that should not be left to evaluators and the testing community as a separate, "stove-piped" activity but one that integrates evaluation with instructional design.
- No lessons—The notion of monolithic instructional modules intended to achieve instructional objectives will also need to evolve if instruction and performance-aiding conversations are to be supported. Objectives may need to be more finely specified by a more comprehensive hierarchical decomposition than called for by current instructional design. As suggested, a capability is needed to treat instruction not as an art or science, but as engineering where specific outcomes, based on detailed knowledge of the learner/user matched with comprehensive representations of the subject matter, can reliably be achieved by all learners—even when the targeted outcomes themselves are modified on the fly.

Other challenges may well occur to the reader. Issues of privacy and security, integration with our current instructional practices, certification at a distance, and the balance between individual learning and the need for social interaction all remain as topics for research, development, and implementation, but they do not seem as peculiar to the ADL vision for distributed learning as these three.

Are These Learning Environments Worthwhile

Hundreds of evaluations have been performed to assess the interactive instructional capabilities incorporated in ADL. As reviewed in more detail by Fletcher (2003), the case, based on empirical data, for using these technology-mediated learning environments may be roughly summarized as the following:

1. Tailoring instruction (education and training) to the needs of individual students has been found to be both an instructional imperative and an economic impossibility.
2. Technology can, in many cases, make this instructional imperative affordable. Under any appreciable student load, it is less expensive to provide instruction with digital technology than to hire a tutor for each student.
3. Technology-based instruction has been found to be more effective than current classroom instructional approaches in many settings across many subject matters.
4. Technology-based instruction is generally less costly than current instructional approaches, especially when many students are to be trained or when instructional objectives concern operating or maintaining costly equipment.
5. Technology-based instruction has been found to decrease the time needed to reach targeted instructional objectives.

Overall, a rule of "thirds" emerges from assessments of computer-based instruction. Findings suggest that use of interactive instructional technologies reduces the cost of instruction by about one-third, and it either reduces time of instruction by about one-third or it increases the amount of skills and knowledge acquired by about one-third. Similar, if not enhanced, results can reasonably be expected with the instructional capabilities that ADL adds to basic technology-based instruction. These results, combined with anywhere, anytime accessibility also provided by technology-based instruction, suggest the value of achieving the ADL vision.

Distance Learning and ADL

ADL approaches contrast with less interactive, less agile, and less flexible technologies such as video teletraining, video conferencing, instructional radio, paper-based correspondence instruction, and instructional telephone, all of which have been used to provide distance-learning. In general, distance learning studies using these less interactive technologies find that they provide instruction that is about as effective as residential classroom instruction, less preferred by students, but notably less costly. For instance, Russell (1999) identified 355 studies reporting no significant differences between distance education and other instructional approaches. His findings are confirmed by other researchers (Bernard, Lou, & Abrami, 2003; Lockee, Burton, & Cross, 1999; Phipps & Merisotis, 1999). Bernard et al. (2004) reported a meta-analytic review of 232 studies comparing distance education with classroom instruction. Among other things, they found superior results for classroom instruction when synchronous approaches to distance education were used and superior results for distance education when asynchronous approaches were used. In a meta-analytic review of 105 empirical studies comparing Web-based learning with classroom approaches, Sitzmann and Wisher (2005) found Web-based instruction to be more effective for teaching declarative knowledge, but they found virtually no difference in the effectiveness of the two forms of instruction for teaching procedural knowledge. In any case, it may be worth emphasizing that lower costs and enhanced accessibility suggest superior cost-effectiveness for distance education even when research finds it to be no more than equally as effective as classroom instruction.

FINAL WORD

Much remains to be done. The vision or view of the future presented in this chapter is not likely to accomplished soon, but it also seems likely, given our progress in such areas as electronics, computer technology, computer communications, and knowledge representation. Serious issues remain in the development of instructional strategies that reliably lead from the learner's (or user's) present state of knowledge, skill, and performance to one that is targeted and desired. We need a capability that is neither art nor science, but instead is most analogous to engineering where known principles are applied to achieved specified outcomes in, if we are fortunate, a cost-effective manner.

The anytime, anywhere objectives of ADL are not contrary to classroom instruction, but very different. They will require changes in roles and responsibilities of students, instructors, and administrators. The budgeting practices and organizational structures now focused on classroom settings will also require major modifications. Like all changes, these are likely to be painful and most certainly difficult to achieve. The prize, however, may be worth it. Enabling the totality of human knowledge to be affordable and available to every individual who seeks it seems a worthy goal. We may well wish that both the technical and administrative difficulties encountered in realizing this vision can and will be surmounted.

REFERENCES

Atkinson, R. C., & Wilson, H. A. (1969). *Computer-assisted instruction; A book of readings.* New York: Academic Press.

Barr, A., Beard, M., & Atkinson, R. C. (1975). A rationale and description of a CAI program to teach the BASIC programming language. *Instructional Science, 4*, 1–31.

Bernard, R. M., Abrami, P. C., Lou, Y., Borokhovski, E., Wade, A., Wozney, L., Wallet, P. A., Fiset, M., & Huang, B. (2004). How does distance education compare with classroom instruction? A meta-analysis of the empirical literature. *Review of Educational Research, 74*, 379–439.

Bernard, R. M., Lou, Y., & Abrami, P. C. (2003). *Is distance education equivalent to classroom instruction? A meta-analysis of the empirical literature.* Paper delivered at the annual convention of the American Educational Research Association, Chicago, IL.

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American, 284*, 34–43.

Brown, J. S., Burton, R. R., & DeKleer, J. (1982). Pedagogical, natural language and knowledge engineering in SOPHIE I, II, and III. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 227–282). New York: Academic Press.

Bryson, J. J., Martin, D. L., McIlraith, S. A., & Stein, L. A. (2002). Toward behavioral intelligence in the semantic web. *Computer, 35*, 48–54.

Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man-Machine Systems, 11*, 190–202.

Chandrasekaran, B., Josephson, J. R., & Benjamins, V. R. (1999). Ontologies: What are they? Why do we need them? *IEEE Intelligent Systems and their applications, 14*, 20–26.

Chipman, S. F. (2003). Gazing yet again into the silicon chip: The future of computers in education. In H. F. O'Neil, Jr. & R. Perez (Eds.), *Technology applications in education: A learning view* (pp. 31–54). Hillsdale, NJ: Lawrence Erlbaum Associates.

Conati, C., Gertner, A., VanLehn, K., & Druzdzel, M. (1997). On-line student modeling for coached problem solving using Bayesian Networks. In A. Jameson, C. Paris, & C. Tasso (Eds.), *Proceedings of UM-97, Sixth International Conference on user modeling.* New York: Springer Wien.

Corbett, A. T., Koedinger, K. R., & Anderson, J. R. (1997). Intelligent tutoring systems. In M. G. Helander, T. K. Landauer, & P. V. Prabhu (Eds.), *Handbook of human-computer interaction* (pp. 849–874). Amsterdam, The Netherlands: Elsevier.

Crowder, N. A. (1959). Automatic teaching by means of intrinsic programming. In E. Galanter (Ed.), *Automatic teaching: The state of the art* (pp. 109–116). New York: Wiley.

Dodds, P. V. W. (Ed.). (2002). *Sharable courseware object reference model (SCORM) version 1.2* (IDA Document D-2677). Alexandria, VA: Institute for Defense Analyses. [Evolving versions of SCORM are posted on line at http://www.idanet.org]

Dodds, P. V. W., & Fletcher, J. D. (2004). Opportunities for new "smart" learning environments enabled by next generation web capabilities. *Journal of Education Multimedia and Hypermedia, 13*(4), 391–404.

Duffy, T. M., & Jonassen, D. H. (Eds.). (1992). *Constructivism and the technology of instruction: A conversation.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Fletcher, J. D. (1975). Models of the learner in computer-assisted instruction. *Journal of Computer-Based Instruction, 3*, 118–126.

Fletcher, J. D. (1988). Intelligent training systems in the military. In S. J. Andriole & G. W. Hopple (Eds.), *Defense applications of artificial intelligence: Progress and prospects* (pp. 33–59). Lexington, MA: Lexington Books.

Fletcher, J. D. (2003). Evidence for learning from technology-assisted instruction. In H. F. O'Neil, Jr. & R. Perez (Eds.), *Technology applications in education: A learning view* (pp. 79–99). Hillsdale, NJ: Lawrence Erlbaum Associates.

Fletcher, J. D., & Johnston, R. (2002). Effectiveness and cost benefits of computer-based aids for maintenance operations. *Computers in Human Behavior, 18*, 717–728.

Fletcher, J. D., & Rockway, M. R. (1986). Computer-based training in the military. In J. A. Ellis (Ed.), *Military contributions to instructional technology* (pp. 171–222). New York: Praeger.

Fletcher, J. D., & Tobias, S. (2003). *Implications of advanced distributed learning for education* (Urban Diversity Series). New York: ERIC Clearinghouse on Urban Education, Teachers College, Columbia University. Retrieved March 17, 2005, from http://iume.tc.columbia.edu/eric_archive/mono/UDS118.pdf

Galanter, E. (Ed.). (1959). *Automatic teaching: The state of the art.* New York: Wiley.

Gibbons, A. S., Nelson, J., & Richards, R. (2000). The nature and origin of instructional objects. In D. Wiley (Ed.), *The instructional use of learning objects*. Retrieved June 22, 2004, from http//:www.reusability.org/read

Graesser, A. C., Gernsbacher, M. A., & Goldman, S. (Eds.). (2003). *Handbook of discourse processes.* Mahwah, NJ: Lawrence Erlbaum Associates.

Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-on-one tutoring. *Applied Cognitive Psychology, 9*, 495–522.

IMS reusable definition of competency or educational objective—information model: version 1.0 final specification. (2002). IMS Global Learning Consortium, Inc. Retrieved September 3, 2002, from http://www.imsglobal.org/competencies/rdceov1p0/imsrdceo_infov1p0.html

James, W. (1890/1950). *Principles of psychology: Volume I.* New York: Dover Press.

Kahn, R., & Wilensky, R. (1995). *A framework for distributed digital object services.* Retrieved March 9, 2005, from http://www.cnri.reston.va.us/k-w.html

Keller, F. S. (1968). Goodbye, teacher .... *Journal of Applied Behavior Analysis, 1*, 79–89.

Keller, F. S. (1985). Lightning strikes twice. *Teaching of Psychology, 12*, 4–8.

Kulik, J. A., Kulik, C-L. C., & Cohen, P. A. (1979). A meta-analysis of outcome studies of Keller's personalized system of instruction. *American Psychologist, 38*, 307–318.

Lockee, B. B., Burton, J. K., & Cross, L. H. (1999). No comparison: Distance education finds a new use for "no significant difference". *Educational Technology Research and Development, 47*, 33–42.

Mann, C. C. (2000). The end of Moore's Law? *Technology Review.* Retrieved December 2, 2004, from http://www.techreview.com/articles/may00/mann.htm

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton, Century, Crofts.

Phipps, R. A., & Merisotis, J. O. (1999). *What's the difference? A review of contemporary research on the effectiveness of distance learning in higher education.* Washington, DC: The Institute of Higher Education Policy.

Psotka, J., Massey, L. D., & Mutter, S. A. (Eds.). (1988). *Intelligent tutoring systems: Lessons learned.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Roschelle, J., & Kaput, J. (1995). Educational software architecture and systemic impact: The promise of component software. *Journal of Educational Computing Research, 14*, 217–228.

Russell, T. L. (1999). *The no significant differences phenomenon.* Chapel Hill, NC: Office of Instructional Telecommunications, North Carolina State University.

SCORM best practices guide for content developers. (2003). Pittsburgh, PA: Learning Systems Architecture Lab, Carnegie Mellon University. Retrieved November 1, 2003, from http://www.lsal.cmu.edu

Service, R. E. (1996). Can chip devices keep shrinking? *Science, 274*, 1834–1836.

Sitzmann, T. M., & Wisher, R. (2005). The effectiveness of Web-based training compared to classroom: A meta-analysis. In S. Carliner & B. Sugrue (Eds.), *ASTD research practice conference proceedings* (pp. 181–187). Alexandria, VA: American Society of Training and Development.

Sleeman, D., & Brown, J. S. (Eds.). (1982). *Intelligent tutoring systems.* New York: Academic Press.

Skinner, B. F. (1954). The science of learning and the art of teaching. *Harvard Education Review, 24*, 86–97.

Soller, A., & Lesgold, A. (2003). A computational approach to analyzing online knowledge sharing interaction. *Proceedings of Artificial Intelligence in Education 2003* (pp. 253–260). Sydney, Australia.

Spohrer, J., Sumner, T., & Shum, S. B. (1998). Educational authoring tools and the educational object economy: Introduction to the special issue from the East/West group. *Journal of Interactive Media in Education.* Retrieved January 20, 2004, from http://www-jime.open.ac.uk/98/10/spohrer-98-10-paper.html

Suppes, P. (Ed.). (1981). *University-level computer assisted instruction at Stanford: 1968–1980.* Stanford, CA: Institute for Mathematical Studies in the Social Sciences.

Suppes, P., & Morningstar, M. (1972). *Computer-assisted instruction at Stanford 1966–1968: Data, models, and evaluation of the arithmetic programs.* New York: Academic Press.

Tobias, S. (1989). Another look at research on the adaptation of instruction to student characteristics. *Educational Psychologist, 24*, 213–227.

Tobias, S., & Frase, L. T. (2000). Educational psychology and training. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 3–24). New York: Macmillan Library Reference.

Towne, D. M. (1998). *Development of scenario tutors in a generalized authoring environment: Feasibility study* (ONR Final Rep. No. 119). Los Angeles: Behavioral Technology Laboratories, University of Southern California.

Towne, D. M. (2003). Automated knowledge acquisition for intelligent support of diagnostic reasoning. In T. Murray, S. Blessing, & S. Ainsworth (Eds.), *Authoring tools for advanced technology learning environments* (pp. 121–147). Norwell, MA: Kluwer Academic Publishers.

Wiley, D. (2000). *The instructional use of learning objects.* Retrieved June 28, 2001, from http://www.reusability.org/read

Wisher, R. A., & Fletcher, J. D. (2004). The case for advanced distributed learning. *Information & Security: An International Journal, 14*, 17–25.

Woolf, B. P., & Regian, J. W. (2000). Knowledge-based training systems and the engineering of instruction. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 339–356). New York: Macmillan Library Reference.